IEEE Instrumentation and Measurement
Technology Conference
Brussels, Belgium, June 4-6, 1996

A Wideband Sampling Voltmeter

T. M. Souders*, B. C. Waltrip, O. B. Laug, and J. P. Deyst
Rm. B162, Bldg. 220
National Institute of Standards and Technology
Gaithersburg, MD USA 20899

*(301) 975-2406 FAX: (301) 926-3972 souders@eeel.nist.gov

Abstract - A high accuracy sampling voltmeter, designed to span the frequency range of 10 Hz to 200 MHz, is described. The instrument operates autonomously, at a measurement update rate of at least one per second. A novel quasi-equivalent time sampling process is used, with a custom strobed comparator as the sampling device and decision element. The architecture and control are presented, along with the time-base design principles. Major error sources associated with the time-base are also discussed.

I. INTRODUCTION

A voltmeter based on sampling principles is being developed at NIST for making accurate rms measurements of repetitive signals ranging in frequency from 10 Hz to 200 MHz. The instrument operates autonomously, without significant operator setup or intervention, and at a measurement update rate of at least one per second over most of its frequency range. In addition to the wide bandwidth capability, the new voltmeter is expected to have very low measurement uncertainties (2×10^{-5} between 10 Hz and 100 kHz, 1×10^{-4} up to 1 MHz, and 2×10^{-3} at 50 MHz). The instrument is intended to provide higher performance, based on bandwidth, measurement accuracy, and speed, than more conventional instruments that use thermal converter technology [1]. The rms value is computed from discrete samples spanning a small integer number of periods of the input waveform. Because the wave shape is preserved in the sampled data, a number of waveform parameters are readily computed in addition to rms value, e.g., average and peak values, and total harmonic distortion.

The sampling is based on a novel quasi-equivalent-time process, using a custom 2.3 GHz bandwidth, strobed comparator [2] as the sampling device and decision element. Quantization is accomplished by using the comparator in a successive approximation feedback loop with an 18-bit digital-to-analog converter (DAC).

This work was supported in part by the Air Force AGMC, under contract number F0000096N1UN1.

A. Prior Art

Earlier work at NIST and elsewhere [2,3,4] has shown equivalent-time sampling comparator systems (SCSs) to be capable of providing wideband, high accuracy rms measurements, although at substantial cost in terms of setup and measurement time. A significant degree of expertise on the part of the operator is also required. The sampling voltage tracker and equivalent-time successive approximation analog-to-digital converters (ADCs) are two such systems.

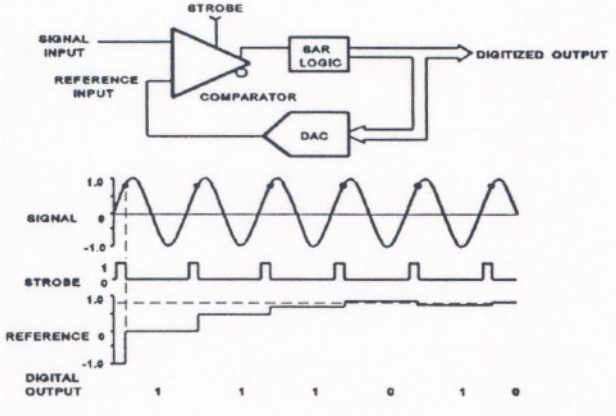

Fig. 1 Equivalent-time successive approximation ADC

The latter system, on which the present work is mostly based, is illustrated in Fig. 1. The principle is similar to that used in a real-time successive approximation ADC; however, rather than using a sample/hold amplifier (S/H) to hold the sampled value during quantization, the signal is resampled (at the same equivalent instant) before each step of the quantization process, by means of the strobed comparator. After one point on the waveform has been completely quantized, the strobing instant is incremented to the next sample time, and the search routine is repeated. This process is continued until a complete record of samples has been obtained. With this architecture, a S/H is not required, and the DAC need not have exceptional dynamic performance since it is

0-7803-3312-8/96/$5.00©1996 IEEE

exercised at a rate far below the equivalent sampling rate. Consequently, the performance is determined almost entirely by the comparator and the time-base.

If the time-base errors are negligible or can be corrected, then only the comparator errors remain. Of these, the most significant for rms measurements (and especially for nearly-sinusoidal signals) is gain flatness. It was shown in [2,3] that the gain flatness of SCSs can be determined very accurately by measuring the step response of the system. Because the step response of an SCS is relatively easy to measure, this becomes a very attractive feature for assessing the overall measurement accuracy, provided the time-base errors are negligible in comparison. The issue of time-base error is addressed in some detail in subsequent sections.

The design and performance of the custom integrated circuit comparator developed for this application was reported previously [2]. A bandwidth of 2.3 GHz was achieved, and the so-called "thermal tails" that compromise the performance of more conventional designs were virtually eliminated.

While current equivalent-time successive approximation systems offer the previously noted advantages, they are not readily adapted to making wideband rms voltage measurements. This is primarily due to the time-base, which has to be carefully set with respect to the input signal frequency, in order to minimize truncation errors (see section V-*A*. below). In addition, in current systems the measurement update rate becomes unacceptably slow as the input signal frequency drops below 10 kHz or so.

B. Objectives

The goals of the current project were to develop an SCS-based voltmeter that could achieve state-of-the-art measurement uncertainties, and operate autonomously, at an acceptable measurement rate that would be nearly independent of signal frequency over most of its frequency range. These goals led to the architecture described in the following sections.

II. DESIGN CONSIDERATIONS

A. Basic Principles and Notation

Consider a sampled sine wave of period T defined by

$$V_m = V_p \sin(m\gamma + \alpha) \qquad (1)$$

where m is the sampling index, $\gamma = 2\pi \Delta t/T$, Δt is the sampling period, and α is the starting phase. The rms value of M samples of V_m is

$$V_{rms} = \left[\frac{1}{M}\sum_{m=0}^{M-1} V_m^2\right]^{1/2} \qquad (2)$$

In the voltmeter, Δt is set so that M samples span an integer number of periods, P, of the input signal. Each sample can be averaged N times, and is quantized with B bits of resolution.

B. Measurement Throughput

The problem of excessive measurement time at low signal frequencies in conventional systems results from time-base designs that only allow for a single sample to be taken per repetition of the input signal, as was illustrated in Fig. 1. To acquire a record of M samples with such a time-base requires a measurement interval spanning at least M periods, although in practice, the situation is much worse: Equivalent-time successive approximation systems only acquire one bit of information per sample, so quantization with B bits of resolution requires B×M periods. Furthermore, if the desired equivalent-time epoch spanned by the record is greater than one signal period, then proportionally longer measurement times are required. Consider a typical example in which 2 periods of a 100 Hz input signal are to be acquired, in a record of 1024 16-bit samples. In this example, the measurement period is at least 328 s. To overcome

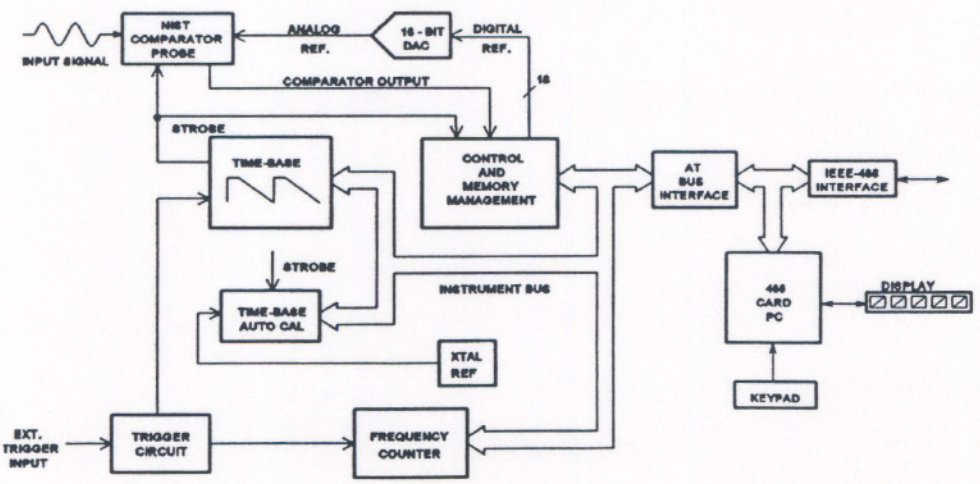

Fig. 2 Block diagram of wideband sampling voltmeter.

this limitation, it is necessary to use a quasi-equivalent-time approach in which multiple samples are taken per repetition of the input signal. In practice, this process is limited only by a minimum duration between samples to allow for setup time and settling. Considering the accuracy goals for the instrument, the minimum practical duration between samples is about 20 μs. Thus for the example case considered above, the measurement period can be reduced to less than 2 s.

III. ARCHITECTURE

Fig. 2 shows a block diagram of the general instrument architecture. As previously indicated, sampling and quantization are accomplished with a strobed comparator circuit operated in an equivalent-time, successive approximation feedback loop with a precision 18-bit DAC.

Sample commands for the comparator are provided by a time-base circuit that employs a conventional ramp technique together with provision to generate multiple sample commands for each repetition of the ramp. Multiple sample commands are generated by causing the reference level (generated by the time-base DAC) to leap ahead of the ramp voltage as soon as the previous level has been crossed. This feature produces subsequent sample commands without waiting for a new ramp to begin. We have dubbed this the "leapfrog" approach. Fig. 3 illustrates this quasi-equivalent-time method. Here is shown an example of eight sample points per period of the waveform that must use two timing ramps to complete the desired number of samples because of the minimum time constraint between samples. The first sample point (0) occurs when the ramp crosses the first DAC level, and subsequent sample points (2, 4, and 6) are generated as the DAC level leaps ahead of the ramp. The remaining sample points are generated in the same way on the next ramp.

The time-base performance is critical to the overall performance of the instrument; consequently, the ramp nonlinearity and timing jitter are tightly controlled over the 7+ decades of time span that it provides, and the ramp slope is accurately maintained via an autocalibration system that performs periodic comparisons to a crystal clock.

Truncation errors in the rms measurement are minimized by monitoring the signal frequency on-the-fly with an integral frequency counter. The instrument then calculates a sampling interval such that a small integer number of cycles of the input signal is exactly spanned with an integer number of samples, and commands the time-base appropriately. (See section V-*A*.)

Global instrument control is provided by a resident single-board 486-type processor. Sampler control and memory management functions, and time-base control are provided by dedicated state machines implemented with programmable logic devices (PLDs). Up to four sampling channels can be controlled simultaneously.

To perform a measurement, the frequency counter measures the frequency of the input signal by counting a corresponding digital output signal from the trigger circuit, and reports it to the processor. The processor then computes and sets the appropriate time-base range, and computes the set of M time-base DAC values that are then loaded into the time-base random access memory (RAM). The triggered time-base then outputs strobe commands in a sequence timed by the RAM contents. For each rms measurement, the M RAM addresses are cycled through B times. Meanwhile, the probe control and memory management board controls the successive approximation process and assembles the B-bit data record from the B×M individual samples. All M most significant bits are acquired first, then the M second most significant bits, and so on. Because of the "leapfrog" process, the samples are acquired in successive groups of modulo-k order rather than in contiguous order; consequently, the memory management PLD reorders them before they are sent to the processor where the rms value (or other parameters of interest) are computed.

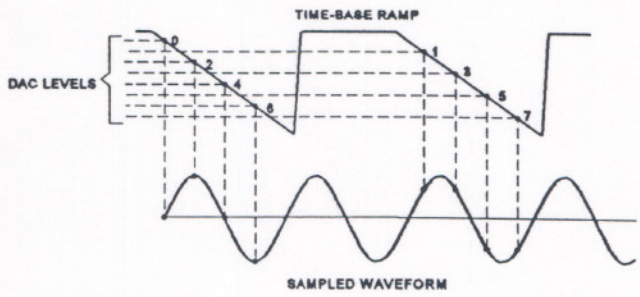

Fig. 3 Quasi-equivalent-time sampling: The "leapfrog" approach.

To achieve the desired accuracy levels at the lower frequency ranges (10 Hz-100 kHz), the random measurement uncertainty caused by timing jitter and comparator noise must be reduced. (The equivalent input noise of the comparator is approximately 400 μV, a consequence of its wide bandwidth.) The most efficient technique, with respect to measurement throughput, is to use the so-called Markov estimator [5] to reduce the noise in each sample. This method requires only N additional repetitions of the time-base ramp cycle to reduce the variance by N, as opposed to B×N repetitions that are required if conventional averaging of N records is used (or B×N repetitions that are required if the record length is increased by a factor of N). The Markov estimator is implemented in the sampler control board by switching the successive approximation control algorithm to the Markov algorithm after the B-bit binary search has taken place. In the Markov algorithm, the DAC input code is incremented or decremented one least significant bit per new sample, depending on the most recent output state of the sampling comparator. This process is repeated N times, and the last N DAC codes are averaged to obtain the Markov estimate.

Time-base autocalibration is accomplished by temporarily applying a reference clock signal to the trigger input and to the

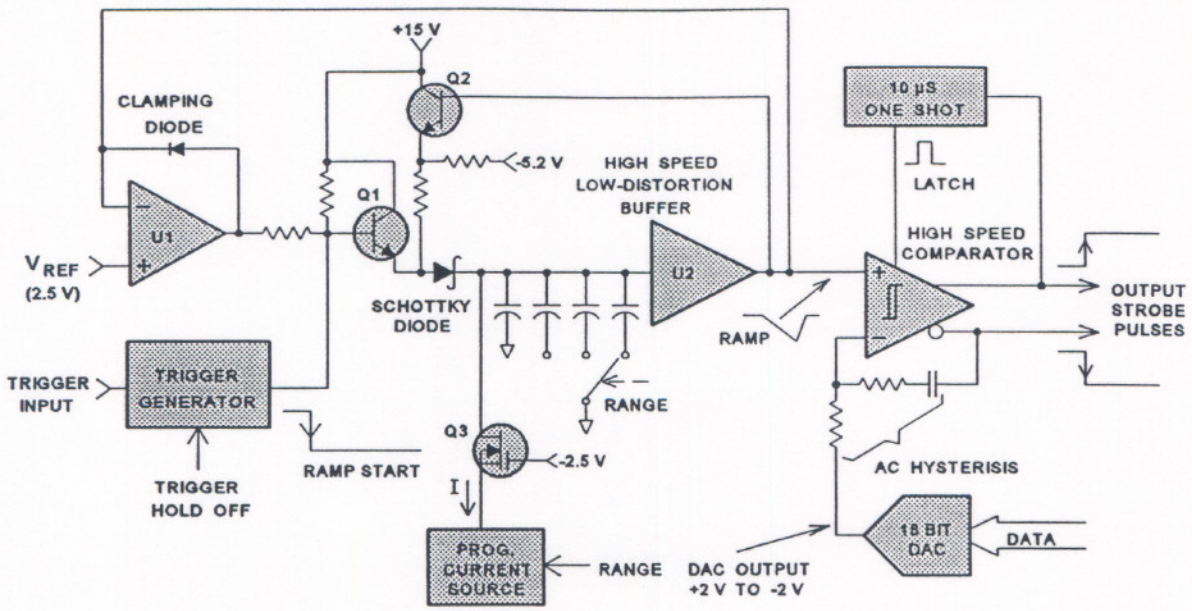

Fig. 4 Simplified diagram of "leapfrog" time-base circuit.

signal input of another strobed comparator. A search is performed for the nth clock transition following the trigger event, by incrementing or decrementing the time-base DAC levels until the transition is found. The comparator's reference input is maintained at mid-transition level. For a positive transition, a clock signal level greater than the reference level causes the time-base to decrement the strobe time, and vise-versa. The time-base RAM values are initially set to position the strobe event in the vicinity of the sought transition, so that the search will be rapid and will converge. After convergence, a Markov estimate is made of the transition location, to minimize the effects of timing jitter. The amounts by which the DAC levels are adjusted are direct measures of the time-base errors. By selecting different values for n, the errors can be determined over the full span of the selected time-base range. Reference clock frequencies of 10 kHz, 1 MHz, and 100 MHz are available, to cover the 20 ranges of the time-base. Both scale-factor and linearity errors can be computed from the measured error data, and subsequent corrections are applied to the RAM values that are calculated during normal instrument operation. While the linearity errors are expected to be stable over time and temperature, the scale factor errors will be temperature dependent, and it is expected that corrections will need to be reevaluated periodically to meet the uncertainty goals at the lower frequencies.

IV. TIME-BASE DESIGN

The time-base circuit produces uniformly-spaced strobe pulses that are synchronized to a trigger signal. These strobe pulses determine the sampling points on the input waveform with respect to the trigger event, and the accuracy of the strobe pulse intervals have a direct bearing on the fidelity of the reconstructed waveform. In addition to maintaining precise timing intervals, the time-base must be able to accommodate the previously described "leapfrog" approach which decreases the acquisition time for low-frequency waveforms.

The time-base design is based on the classical ramp or sweep generator technique in which a trigger signal initiates a linear voltage ramp. The ramp is generated by allowing a constant current to charge a capacitor. A strobe pulse is produced by a voltage comparator at the instant the ramp voltage crosses a preset reference level.

Fig. 4 shows a simplified diagram of the time-base circuit that highlights some of its features which enhance the performance and enable the "leapfrog" technique to be realized. Before the ramp is started, the selected capacitor is maintained at a reference level of 2.5 V in a feedback loop around U1 and U2. When the trigger generator produces a "ramp start" pulse the feedback path is opened by turning Q1 off, which back-biases the Schottky diode and thus allows the programmable current source to charge the selected capacitor. Amplifier U1 is prevented from saturating in the open-loop condition by a clamping diode. The integration of the constant current by the capacitor will produce a linear change of voltage vs. time that descends from the initial starting level of 2.5 V to about -2.5 V. (The portion of the ramp from 2.0 V to -2.0 V is used as the time-base.) At that point a reset circuit (not shown) will reset the trigger generator so that Q1 is turned on and the capacitor is charged to the 2.5 V level again. After a suitable hold-off time to allow for the voltage to settle across the capacitor, the next available trigger will restart the ramp. A strobe pulse is produced by the comparator each time the ramp voltage crosses the preset DAC output voltage.

The performance of the programmable current source is enhanced by a common-gate N-channel MOS FET (Q3). This

stage raises the equivalent output impedance of the current source to improve the ramp linearity while adding only about 1 pF of parasitic capacitance to the integrating node. More importantly, this parasitic capacitance tends to remain reasonably constant over the voltage range of the ramp, thereby maintaining good linearity for the smallest value of integrating capacitance.

A secondary feedback loop from the output of buffer U2 through emitter follower stage Q2 to the anode of the Schottky diode operates during the ramp period by maintaining a constant back-bias of about 0.8 V on the Schottky diode. This technique practically eliminates the voltage dependent effect of the diode capacitance from contributing to ramp nonlinearity.

A buffer amplifier (U2) between the integrating capacitor and the comparator is required for the "leapfrog" mode of operation. The buffer prevents comparator "kickout" currents from disturbing the ramp between successive samples. The performance requirements of the buffer amplifier are quite demanding and include low distortion, wide bandwidth, and high input impedance. A FET buffer with a cascode input configuration was found to satisfy the needs for this application.

The comparator that produces the output strobe pulses requires some hysteresis to minimize the likelihood of retriggering caused by noise, and a timed latch signal to prevent retriggering when the DAC is updated. The ac voltage hysteresis loop is designed to provide a large amount of almost instantaneous positive feedback with a short time constant to help reduce noise-induced retriggering. The time constant must be sufficiently short so that the effects of the loop die out rapidly after each comparator transition. The slower secondary loop with a 10 μs period latches the state of the comparator to prevent false strobe pulses while the DAC is reprogrammed and settles to the new level. When the latch is released the comparator is reset because the DAC output level is now below (or ahead of) the descending ramp. When the ramp level reaches the new DAC output level another strobe pulse is generated. The process repeats itself along the ramp at a minimum time interval of 20 μs between samples.

The time-base is designed for up to 20 ranges in a 1-2-5 sequence by selecting one of four capacitors and one of six current-source levels. The smallest value of capacitance is always in the circuit while the larger values are switched in as needed. The ramp spans range from 100 ns to 250 ms.

The time-base linearity, measured over the full span of each range, is within 100 ppm on all but the 100 ns range. For this range, the linearity degrades to about 300 ppm. The 1-σ strobe jitter is <200 ppm of the full-scale span for all ranges.

V. ERROR SOURCES

A. Truncation Errors

Truncation errors occur when a non-integer number of periods of the input signal are sampled; consequently, the sample times must be computed and set very accurately. To assure the lowest truncation errors, the exact scale-factor of the time-base range is automatically measured periodically via the time-base autocalibration circuit described in section III. As shown in [6], the truncation error (assuming it is <<1), relative to the true *mean squared* value $V_P^2/2$, can be represented as

$$\frac{1}{M}\cdot\frac{\sin\delta}{\sin\gamma}\cos(2\alpha-\delta-\gamma) \qquad (3)$$

where the truncation angle δ is the difference between the summation interval, given by $M\gamma$, and the nearest integral number of periods, P, of the sine wave, i.e., $\delta \equiv 2\pi P - M\gamma$. Expressing the sampling period Δt as $(T/M)(1+\mu)$, where μ is the proportional error in setting the time-base scale factor with respect to the signal period, and assuming that δ, γ and (3) are all <<1, the truncation error in the *rms value* can be approximated as

$$\frac{\mu}{2}\cos(2\alpha-\delta-\gamma) \le \mu/2 \qquad (4)$$

If the time-base scale factor is perfect, then $M\Delta t$ spans an exact integer number of periods, and the truncation error is zero, regardless of the other terms. As noted in section IV, in the new time-base design, μ is on the order of 100 ppm or less over most of the ranges. Since $\delta/\gamma \approx M\mu$, δ can generally be disregarded in (4) (being $<<\gamma$) for most reasonable choices of M (e.g., ≤ 1024). For any arbitrary angle α, the truncation error should then be no greater than 50 ppm. If less uncertainty is required, then $2\alpha-\gamma$ can be set to approximately $\pi/2$ rad, thus reducing the uncertainty to a negligible level. (For sinusoids, this is accomplished by starting the summation (in (2)) one sample prior to a point where $V_m = V_p/\sqrt{2}$. Note that the limiting error ($\mu/2$) is independent of harmonic order, so long as the small-number approximations are still valid for δ and γ. These approximations hold to ≤ 1% for at least 20 harmonics, with P=1, M=512 and μ=100 ppm.

B. Noise and Jitter

The systematic error in rms value due to additive noise is only significant, i.e., ≥ 5 ppm, when the signal-to-noise ratio is less than about 3×10^{-3}. For a noise level of 400 μV, this implies a signal level of at least 133 mV. Smaller signals will require some averaging to keep the errors at negligible levels. More significant though is the influence of noise on the standard deviation of the rms estimates. For signal amplitudes >> σ_{noise}, and no signal averaging,

$$SDEV[V_{rms}] \approx \frac{\sigma_{noise}}{\sqrt{M}} \qquad (5)$$

For σ_{noise} = 400 μV and M=512, this gives a standard deviation of 18 μV, which is marginal performance even for 1 V signals.

Therefore, signal averaging is even more important for reducing the measurement variance.

In the case of timing jitter, the expectation of the rms estimate is only weakly affected by timing jitter [5]. For example, with $\sigma_t = 10^{-4} \times T_{RAMP}$ and P=1, the rms value is reduced by only 0.2 ppm. As with noise, however, the resulting standard deviation of the estimate (normalized by the expected rms value) is much greater:

$$\frac{SDEV\,[V_{rms}]}{V_p/\sqrt{2}} \approx \sqrt{\frac{2}{M}} \cdot \frac{\pi\,\sigma_t P}{T_{RAMP}} \qquad (6)$$

With $\sigma_t = 10^{-4} \times T_{RAMP}$, P=1 and M=512, this gives a normalized standard deviation of 20 ppm, before averaging. Regarding accumulated jitter [7], in the present time-base design the jitter variance had been found to be essentially the same at the end as at the beginning of the time-base span. This is an indication that the time-base comparator noise, rather than the current source noise, is the dominant source of jitter.

C. Time-Base Quantization and Linearity Errors

Quantization of the discrete time steps causes time-base noise that is similar to jitter; however, its effect cannot be reduced by averaging because it is deterministic and is correlated with the signal. The relation of quantization error to sampling index is dependent on the resolution of the time-base DAC as well as on M, P, and the signal frequency, because all of these are used to compute the time-base range and RAM contents. The errors exhibit non-harmonically related patterns vs. sampling index, m. In pathological cases, the resulting error in the estimate of the rms value of a sine wave can approach the maximum possible error caused by time-base nonlinearity [3], namely $4 \cdot P \cdot \tau$, where τ, the peak proportional time-base error, is in this case $2^{-(b+1)}$, with b the resolution of the time-base DAC, in bits. For this reason, an 18-bit DAC was chosen for use in the time-base, guaranteeing errors in the rms estimates of less than 8 ppm (for P=1).

The maximum possible error referred to above, $4 \cdot P \cdot \tau$, is unrealistically large for most sources of time-base linearity error that are encountered. For example, a typical source of linearity error in the time-base is that caused by a leakage path shunting the integrating capacitor. This produces a bow in the voltage ramp given by

$$V(t) = IR\,(1 - e^{-t/RC}) \qquad (7)$$

where I is the constant current, R is the shunt resistance, and C is the integrating capacitance. This type of error produces a maximum proportional error in the rms estimate (of a single period sine wave) of 0.16 τ (vs. 4 τ), where τ is the peak linearity error relative to the duration of the ramp, T_{RAMP}.

VI. CONCLUSION

The basic sampling voltmeter described above is currently operational, but some key features, notably the frequency counter and the time-base autocalibration circuits, have yet to be fully implemented. The expected measurement uncertainties are based on earlier measurements obtained using the NIST sampling comparator in a conventional SCS configuration.

ACKNOWLEDGEMENT

Robert Palm of NIST performed innumerable drafting, layout, and assembly services during the course of this development. His talents and contributions are greatly appreciated.

REFERENCES

[1] L. Eccleston, "Designing an ac measurement standard," Proceedings 1991 Meas. Science Conf.

[2] O. B. Laug, T. M. Souders and D. R. Flach, "A custom integrated circuit comparator for high-performance sampling applications," IEEE Trans. Instrum. Meas., vol. 41, pp 850-855, December 1992.

[3] T. M. Souders and P. S. Hetrick, "Accurate rf voltage measurements using a sampling voltage tracker," IEEE Trans. Instrum. Meas., vol. IM-38, pp 451-456, April 1989.

[4] C. Gyles, "Repetitive waveform high frequency, high resolution digitizer," IEEE Trans. Instrum. Meas., vol;. 38, pp. 917-920, August 1989.

[5] T. M. Souders, et al, "Effects of timing jitter in sampling systems," IEEE Trans. Instrum. Meas., vol.39, PP. 80-85, February, 1990.

[6] G. N. Stenbakken, "A wideband sampling wattmeter,"IEEE Trans. Power App. and Sys., vol. 103, October 1984.

[7] S. S. Awad, "The effects of accumulated timing jitter on some sine wave measurements," IEEE Trans. Instrum. Meas., vol. 44, p 945, October 1995.